



Welcome to IEEE Xplore®

- ☐ Home
- ☐ What Can I Access?
- ☐ Log-out

Tables of Contents

- ☐ Journals & Magazines
- ☐ Conference Proceedings
- ☐ Standards

Search

- ☐ By Author
- ☐ Basic
- ☐ Advanced
- ☐ CrossRef

Member Services

- ☐ Join IEEE
- ☐ Establish IEEE Web Account
- ☐ Access the IEEE Member Digital Library

IEEE Enterprise

- ☐ Access the IEEE Enterprise File Cabinet

Your search matched **0** of **1128145** documents.

A maximum of **500** results are displayed, **15** to a page, sorted by **Relevance Descending** order.

Refine This Search:

You may refine your search by editing the current search expression or entering a new one in the text box.

☐ Check to search within this result set

Results Key:

JNL = Journal or Magazine **CNF** = Conference **STD** = Standard

Results:

No documents matched your query.


[Print Format](#)

[Home](#) | [Log-out](#) | [Journals](#) | [Conference Proceedings](#) | [Standards](#) | [Search by Author](#) | [Basic Search](#) | [Advanced Search](#) | [Join IEEE](#) | [Web Account](#) | [New this week](#) | [OPAC Linking Information](#) | [Your Feedback](#) | [Technical Support](#) | [Email Alerting](#) | [No Robots Please](#) | [Release Notes](#) | [IEEE Online Publications](#) | [Help](#) | [FAQ](#) | [Terms](#) | [Back to Top](#)



Terms used

pinned or **pinning** **near/5** **partition** or **sector** or **section** or **divi** or **part** or **block** or **bound** and **pinned** or **pinning**

Sort results by

Display results
☒ Save results to a Binder

☒ Search Tips

☐ Open results in a new window

Try an [Advanced Search](#)

Try this search in [The A](#)

Results 1 - 20 of 200

Best 200 shown

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

1 [Cellular disco: resource management using virtual clusters on shared-memory multiprocessors](#)

Kinshuk Govil, Dan Teodosiu, Yongqiang Huang, Mendel Rosenblum

August 2000 **ACM Transactions on Computer Systems (TOCS)**, Volume 18 Issue 3

Full text available: ☒ pdf(287.05 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index t](#)

Despite the fact that large-scale shared-memory multiprocessors have been commercially available, system software that fully utilizes all their features is still not available, mostly due to the complexity of making the required changes to the operating system. A recently proposed approach, called Disco, reduces this development cost by using a virtual machine monitor that leverages the existing operating system technology. In this paper we present a ...

Keywords: fault containment, resource management, scalable multiprocessors, virtual machines

2 [Cache Memories](#)

Alan Jay Smith

September 1982 **ACM Computing Surveys (CSUR)**, Volume 14 Issue 3

Full text available: ☒ pdf(4.61 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

3 [The Vesta parallel file system](#)

Peter F. Corbett, Dror G. Feitelson

August 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 3

Full text available: ☒ pdf(649.08 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index t](#)

The Vesta parallel file system is designed to provide parallel file access to application programs running on multicomputers with parallel I/O subsystems. Vesta uses a new abstraction of files: a file is not a sequence of bytes but rather it can be partitioned into multiple disjoint sequences that are accessed in parallel. The partitioning can also be changed dynamically—reduces the need for synchronization and coordination during the execution of control over the layout ...

Keywords: data partitioning, parallel computing, parallel file system

4 [Virtual machine monitors: Xen and the art of virtualization](#)

Paul Barham, Boris Dragovic, Keir Fraser, Steven Hand, Tim Harris, Alex Ho, Rolf Neugebauer, Ian Pratt, Paul Warfield

October 2003 **Proceedings of the nineteenth ACM symposium on Operating systems principles**

Full text available: ☒ pdf(168.76 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index t](#)

Numerous systems have been designed which use virtualization to subdivide the ample resources

computer. Some require specialized hardware, or cannot support commodity operating systems. Some sacrifice binary compatibility at the expense of performance. Others sacrifice security or functionality for speed of resource isolation or performance guarantees; most provide only best-effort provisioning, risking correctness of service. This paper presents Xen, an x86 virtual machine monitor ...

Keywords: hypervisors, paravirtualization, virtual machine monitors

5 Cellular Disco: resource management using virtual clusters on shared-memory multiprocessors

Kinshuk Govil, Dan Teodosiu, Yongqiang Huang, Mendel Rosenblum

December 1999 **ACM SIGOPS Operating Systems Review , Proceedings of the seventeenth ACM SIGOPS Operating systems principles**, Volume 33 Issue 5

Full text available: [pdf\(1.93 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Despite the fact that large-scale shared-memory multiprocessors have been commercially available, system software that fully utilizes all their features is still not available, mostly due to the complexity of making the required changes to the operating system. A recently proposed approach, called Disco, reduces this development cost by using a virtual machine monitor that leverages the existing operating system technology. In this paper we present a system ...

6 Static single assignment form for machine code

Allen Leung, Lal George

May 1999 **ACM SIGPLAN Notices , Proceedings of the ACM SIGPLAN 1999 conference on programming language design and implementation**, Volume 34 Issue 5

Full text available: [pdf\(1.31 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Static Single Assignment (SSA) is an effective intermediate representation in optimizing compilers. Traditional SSA form and optimizations are not applicable to programs represented as native machine code because the use of dedicated registers imposed by calling conventions, the runtime system, and the hardware must be made explicit. We present a simple scheme for converting between programs in machine code such that references to dedicated physical registers ...

7 An architecture for secure wide-area service discovery

Todd D. Hodes, Steven E. Czerwinski, Ben Y. Zhao, Anthony D. Joseph, Randy H. Katz

March 2002 **Wireless Networks**, Volume 8 Issue 2/3

Full text available: [pdf\(365.68 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The widespread deployment of inexpensive communications technology, computational resources, and network-enabled end devices poses an interesting problem for end users: how to find a particular network service or device out of hundreds of thousands of accessible services and devices. This paper presents the architecture and implementation of a secure wide-area Service Discovery Service (SDS). SDS providers use the SDS to advertise descriptions of available ...

Keywords: location services, name lookup, network protocols, service discovery

8 The Alpine file system

M. R. Brown, K. N. Koling, E. A. Taft

November 1985 **ACM Transactions on Computer Systems (TOCS)**, Volume 3 Issue 4

Full text available: [pdf\(2.95 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Alpine is a file system that supports atomic transactions and is designed to operate as a service over a network. Alpine's primary purpose is to store files that represent databases. An important secondary purpose is to store ordinary files representing documents, program modules, and the like. Unlike other file servers designed for scientific literature, Alpine uses a log-based technique to implement atomic file update. Another unusual aspect is that it performs all communication ...

9 Distributed transactions for reliable systems

Alfred Z. Spector, Dean Daniels, Daniel Duchamp, Jeffrey L. Eppinger, Randy Pausch

December 1985 **ACM SIGOPS Operating Systems Review , Proceedings of the tenth ACM symposium on Operating systems principles**, Volume 19 Issue 5

Full text available: [pdf\(1.44 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

10 Secure buffering in firm real-time database systems

Binto George, Jayant R. Haritsa

February 2000 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 17

Full text available:  pdf(227.42 KB)

Additional Information: [full citation](#), [abstract](#), [index terms](#)


Many real-time database applications arise in electronic financial services, safety-critical installation systems where enforcing is crucial to the success of the enterprise. We investigate here the performance in terms of killed transactions, of guaranteeing *multi-level secrecy* in a real-time database system applications with *firm* deadlines. In particular, we focus on the *buffer management* aspects of this contributions a ...

Keywords: Buffer management, Covert channels, Firm deadlines, Real-time database

11 Query evaluation techniques for large databases

Goetz Graefe

June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

Full text available:  pdf(9.37 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. Traditional query-oriented and extensible database systems will not solve this problem. On the contrary, modern database systems exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's systems manipulate simple records, query-processing algorithms must be able to handle complex objects ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, object-oriented database systems, operator model of parallelization, parallel algorithms, relational set-matching algorithms, sort-hash duality

12 I/O: miNI: reducing network interface memory requirements with dynamic handle lookup

Reza Azimi, Angelos Bilas

June 2003 **Proceedings of the 17th annual international conference on Supercomputing**

Full text available:  pdf(289.75 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)


Recent work in low-latency, high-bandwidth communication systems has resulted in building user-space Network Interface Controllers (NICs) and communication abstractions that support direct access from the NIC to virtual memory to avoid both data copies and operating system intervention. Such mechanisms require the NIC to directly manipulate user-space communication buffers for delivering data and achieving protection. To achieve these abilities, NICs must maintain appropriate tables ...

Keywords: parallel architectures, system area networks

13 Virtual machines: Memory resource management in VMware ESX server

Carl A. Waldspurger

December 2002 **ACM SIGOPS Operating Systems Review**, Volume 36 Issue SI

Full text available:  pdf(1.65 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

VMware ESX Server is a thin software layer designed to multiplex hardware resources efficiently among virtual machines running unmodified commodity operating systems. This paper introduces several novel mechanisms and policies for managing memory. A *ballooning* technique reclaims the pages considered free by the operating system running in a virtual machine. An *idle memory tax* achieves efficient memory management while maintaining performance isolation guarantees ...

14 Programming languages: Compiler-assisted demand paging for embedded systems with flash

Chanik Park, Junghee Lim, Kiwon Kwon, Jaejin Lee, Sang Lyul Min

September 2004 **Proceedings of the fourth ACM international conference on Embedded software**

Full text available:  pdf(392.66 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this paper, we propose a novel, application specific demand paging mechanism for low-end emul with flash memory as secondary storage. These systems are not equipped with virtual memory. A space called an execution buffer is allocated to page an application. An application-specific page in the buffer. The manager is generated by a compiler post-pass and combined with the application in post-pass analyzes the ELF executable image of an appl ...

Keywords: SRAM, clustering, compilers, embedded systems, flash memory, heterogeneous memory pass optimization

15 Physical Experimentation with Prefetching Helper Threads on Intel's Hyper-Threaded Processors

Dongkeun Kim, Steve Shih-wei Liao, Perry H. Wang, Juan del Cuvillo, Xinmin Tian, Xiang Zou, Hong Yeung, Milind Girkar, John P. Shen

March 2004 **Proceedings of the international symposium on Code generation and optimization directed and runtime optimization**

Full text available:  pdf(264.47 KB)


Additional Information: [full citation](#), [abstract](#), [citations](#)

Pre-execution techniques have received much attention as an effective way of prefetching cache blocks to reduce ever-increasing memory latency. A number of pre-execution techniques based on hardware, compiler, and OS have been proposed and studied extensively by researchers. They report promising results on simulators and on Simultaneous Multithreading (SMT) processor. In this paper, we apply the helper threading idea on multithreaded machine, i.e., Intel Pentium 4 processor with Hyper-Threading ...

16 Cluster communication protocols for parallel-programming systems

Kees Verstoep, Raoul A. F. Bhoedjang, Tim Rühl, Henri E. Bal, Rutger F. H. Hofman

August 2004 **ACM Transactions on Computer Systems (TOCS)**, Volume 22 Issue 3

Full text available:  pdf(1.29 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Clusters of workstations are a popular platform for high-performance computing. For many parallel applications, efficient use of a fast interconnection network is essential for good performance. Several modern networks include programmable network interfaces that can be tailored to perform protocol tasks that would need to be done by the host processors. Finding the right trade-off between protocol processing and the network interface is difficult in general. In ...

Keywords: Clusters, parallel-programming systems, system area networks

17 The case for SRAM main memory

Philip Machanick

December 1996 **ACM SIGARCH Computer Architecture News**, Volume 24 Issue 5

Full text available:  pdf(549.25 KB)


Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

The growing CPU-memory gap is resulting in increasingly large cache sizes. As cache sizes increase, the benefit of cache becomes less of a win. At the same time, since costs of going to DRAM increase, it becomes more to pin critical data in the cache---a problem if a cache is direct-mapped or has a low degree of associativity. Something else which is a problem for caches of low associativity is reducing misses by using a better cache policy. This paper proposes that L2 cache s ...

18 Resource partitioning in general purpose operating systems: experimental results in Windows NT

D. G. Waddington, D. Hutchison

October 1999 **ACM SIGOPS Operating Systems Review**, Volume 33 Issue 4

Full text available:  pdf(1.56 MB)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

The principal role of the operating system is that of resource management. Its task is to present a set of services to the applications and users it supports. Traditionally, general-purpose operating systems like Windows NT, federate resource sharing in a fair manner, with the predominant goal of efficient resource sharing. As a result the chosen scheduling algorithms are not suited to applications that have stringent Quality of Service and resource management require ...


19 Application performance and flexibility on exokernel systems

M. Frans Kaashoek, Dawson R. Engler, Gregory R. Ganger, Héctor M. Briceño, Russell Hunt, David M. Pinckney, Robert Grimm, John Jannotti, Kenneth Mackenzie

20 Load-sensitive routing of long-lived IP flows

Anees Shaikh, Jennifer Rexford, Kang G. Shin

August 1999 **ACM SIGCOMM Computer Communication Review , Proceedings of the conference**
Applications, technologies, architectures, and protocols for computer commun
Issue 4

Full text available:  pdf(1.57 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index t](#)

Internet service providers face a daunting challenge in provisioning network resources, due to the Internet and wide fluctuations in the underlying traffic patterns. The ability of dynamic routing to congested links and improve application performance makes it a valuable traffic engineering tool. deployment of load-sensitive routing is hampered by the overheads imposed by link-state update selection, and signaling. Under reasonable protoc ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, I
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)



Terms used

pinned or **pinning** **near/5** **partition** or **sector** or **section** or **divi** or **part** or **block** or **bound** and **pinned** or **pinning**

Sort results by

Display results

[Save results to a Binder](#)

[Search Tips](#)
☐ Open results in a new window

Try an [Advanced Search](#)

Try this search in [The A](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

1 [Cellular disco: resource management using virtual clusters on shared-memory multiprocessors](#)

Kinshuk Govil, Dan Teodosiu, Yongqiang Huang, Mendel Rosenblum

August 2000 **ACM Transactions on Computer Systems (TOCS)**, Volume 18 Issue 3

Full text available: [pdf\(287.05 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index t](#)

Despite the fact that large-scale shared-memory multiprocessors have been commercially available, system software that fully utilizes all their features is still not available, mostly due to the complexity making the required changes to the operating system. A recently proposed approach, called Disco, reduces this development cost by using a virtual machine monitor that leverages the existing operating technology. In this paper we present a ...

Keywords: fault containment, resource management, scalable multiprocessors, virtual machines

2 [Cache Memories](#)

Alan Jay Smith

September 1982 **ACM Computing Surveys (CSUR)**, Volume 14 Issue 3

Full text available: [pdf\(4.61 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

3 [The Vesta parallel file system](#)

Peter F. Corbett, Dror G. Feitelson

August 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 3

Full text available: [pdf\(649.08 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index t](#)

The Vesta parallel file system is designed to provide parallel file access to application programs running on multicomputers with parallel I/O subsystems. Vesta uses a new abstraction of files: a file is not a sequence of bytes but rather it can be partitioned into multiple disjoint sequences that are accessed in parallel. The partitioning can also be changed dynamically—reduces the need for synchronization and coordination during the control over the layout ...

Keywords: data partitioning, parallel computing, parallel file system

4 [Virtual machine monitors: Xen and the art of virtualization](#)

Paul Barham, Boris Dragovic, Keir Fraser, Steven Hand, Tim Harris, Alex Ho, Rolf Neugebauer, Ian Pratt, Warfield

October 2003 **Proceedings of the nineteenth ACM symposium on Operating systems principles**

Full text available: [pdf\(168.76 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index t](#)

Numerous systems have been designed which use virtualization to subdivide the ample resources


computer. Some require specialized hardware, or cannot support commodity operating systems. Some lack binary compatibility at the expense of performance. Others sacrifice security or functionality for speed, resource isolation or performance guarantees; most provide only best-effort provisioning, risking correctness of service. This paper presents Xen, an x86 virtual machine monitor ...

Keywords: hypervisors, paravirtualization, virtual machine monitors

5 Cellular Disco: resource management using virtual clusters on shared-memory multiprocessors

Kinshuk Govil, Dan Teodosiu, Yongqiang Huang, Mendel Rosenblum

December 1999 **ACM SIGOPS Operating Systems Review , Proceedings of the seventeenth ACM SIGOPS Operating systems principles**, Volume 33 Issue 5

Full text available:  pdf(1.93 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Despite the fact that large-scale shared-memory multiprocessors have been commercially available, system software that fully utilizes all their features is still not available, mostly due to the complexity of making the required changes to the operating system. A recently proposed approach, called Disco, reduces this development cost by using a virtual machine monitor that leverages the existing operating system technology. In this paper we present a system ...

6 Static single assignment form for machine code

Allen Leung, Lal George

May 1999 **ACM SIGPLAN Notices , Proceedings of the ACM SIGPLAN 1999 conference on programming language design and implementation**, Volume 34 Issue 5

Full text available:  pdf(1.31 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Static Single Assignment (SSA) is an effective intermediate representation in optimizing compilers. Traditional SSA form and optimizations are not applicable to programs represented as native machine code because the use of dedicated registers imposed by calling conventions, the runtime system, and the hardware must be made explicit. We present a simple scheme for converting between programs in machine code such that references to dedicated physical registers ...

7 An architecture for secure wide-area service discovery

Todd D. Hodes, Steven E. Czerwinski, Ben Y. Zhao, Anthony D. Joseph, Randy H. Katz

March 2002 **Wireless Networks**, Volume 8 Issue 2/3

Full text available:  pdf(365.68 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The widespread deployment of inexpensive communications technology, computational resources, and network-enabled end devices poses an interesting problem for end users: how to find a particular network service or device out of hundreds of thousands of accessible services and devices. This paper presents the architecture and implementation of a secure wide-area Service Discovery Service (SDS). SDS providers use the SDS to advertise descriptions of available ...

Keywords: location services, name lookup, network protocols, service discovery

8 The Alpine file system

M. R. Brown, K. N. Kolling, E. A. Taft

November 1985 **ACM Transactions on Computer Systems (TOCS)**, Volume 3 Issue 4

Full text available:  pdf(2.95 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Alpine is a file system that supports atomic transactions and is designed to operate as a service over a network. Alpine's primary purpose is to store files that represent databases. An important secondary purpose is to store ordinary files representing documents, program modules, and the like. Unlike other file servers designed for scientific literature, Alpine uses a log-based technique to implement atomic file update. Another unusual aspect is that it performs all communication over a single network ...

9 Distributed transactions for reliable systems

Alfred Z. Spector, Dean Daniels, Daniel Duchamp, Jeffrey L. Eppinger, Randy Pausch

December 1985 **ACM SIGOPS Operating Systems Review , Proceedings of the tenth ACM symposium on Operating systems principles**, Volume 19 Issue 5

Full text available:  pdf(1.44 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

10 Secure buffering in firm real-time database systems

Binto George, Jayant R. Haritsa

February 2000 **The VLDB Journal – The International Journal on Very Large Data Bases**, Volume 17

Full text available:  pdf(227.42 KB)

Additional Information: [full citation](#), [abstract](#), [index terms](#)


Many real-time database applications arise in electronic financial services, safety-critical installation systems where enforcing is crucial to the success of the enterprise. We investigate here the performance in terms of killed transactions, of guaranteeing *multi-level secrecy* in a real-time database system applications with *firm* deadlines. In particular, we focus on the *buffer management* aspects of this contributions a ...

Keywords: Buffer management, Covert channels, Firm deadlines, Real-time database

11 Query evaluation techniques for large databases

Goetz Graefe

June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

Full text available:  pdf(9.37 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Database management systems will continue to manage large data volumes. Thus, efficient algorithms and manipulating large sets and sequences will be required to provide acceptable performance. The oriented and extensible database systems will not solve this problem. On the contrary, modern database systems exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's systems manipulate simple records, query-processing ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database object-oriented database systems, operator model of parallelization, parallel algorithms, relational set-matching algorithms, sort-hash duality

12 I/O: miNI: reducing network interface memory requirements with dynamic handle lookup

Reza Azimi, Angelos Bilas

June 2003 **Proceedings of the 17th annual international conference on Supercomputing**

Full text available:  pdf(289.75 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)


Recent work in low-latency, high-bandwidth communication systems has resulted in building user-Interface Controllers (NICs) and communication abstractions that support direct access from the NIC virtual memory to avoid both data copies and operating system intervention. Such mechanisms directly manipulate user-level communication buffers for delivering data and achieving protection abilities, NICs must maintain appropriate ...

Keywords: parallel architectures, system area networks

13 Virtual machines: Memory resource management in VMware ESX server

Carl A. Waldspurger

December 2002 **ACM SIGOPS Operating Systems Review**, Volume 36 Issue SI

Full text available:  pdf(1.65 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

VMware ESX Server is a thin software layer designed to multiplex hardware resources efficiently across virtual machines running unmodified commodity operating systems. This paper introduces several novel mechanisms and policies for managing memory. A *ballooning* technique reclaims the pages considered by the operating system running in a virtual machine. An *idle memory tax* achieves efficient memory maintaining performance isolation guarantee ...

14 Programming languages: Compiler-assisted demand paging for embedded systems with flash

Chanik Park, Junghee Lim, Kiwon Kwon, Jaejin Lee, Sang Lyul Min

September 2004 **Proceedings of the fourth ACM international conference on Embedded software**

Full text available:  pdf(392.66 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this paper, we propose a novel, application specific demand paging mechanism for low-end emul with flash memory as secondary storage. These systems are not equipped with virtual memory. A space called an execution buffer is allocated to page an application. An application-specific page r the buffer. The manager is generated by a compiler post-pass and combined with the application i post-pass analyzes the ELF executable image of an appl ...

Keywords: SRAM, clustering, compilers, embedded systems, flash memory, heterogeneous mem pass optimization

15 Physical Experimentation with Prefetching Helper Threads on Intel's Hyper-Threaded Proce:

Dongkeun Kim, Steve Shih-wei Liao, Perry H. Wang, Juan del Cuvillo, Xinmin Tian, Xiang Zou, Hong Yeung, Milind Girkar, John P. Shen

March 2004 **Proceedings of the international symposium on Code generation and optimizat directed and runtime optimization**

Full text available:  pdf(264.47 KB)


Additional Information: [full citation](#), [abstract](#), [citations](#)

Pre-execution techniques have received much attention as an effective way of prefetching cache bl ever-increasing memory latency. A number of pre-execution techniques based on hardware, compil been proposed and studied extensively by researchers. They report promising results on simulators Simultaneous Multithreading (SMT) processor. In this paper, we apply the helper threading idea on multithreaded machine, i.e., Intel Pentium 4 processor with Hyp ...

16 Cluster communication protocols for parallel-programming systems

Kees Verstoep, Raoul A. F. Bhoedjang, Tim Rühl, Henri E. Bal, Rutger F. H. Hofman

August 2004 **ACM Transactions on Computer Systems (TOCS)**, Volume 22 Issue 3

Full text available:  pdf(1.29 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Clusters of workstations are a popular platform for high-performance computing. For many paralle efficient use of a fast interconnection network is essential for good performance. Several modern Networks include programmable network interfaces that can be tailored to perform protocol tasks would need to be done by the host processors. Finding the right trade-off between protocol proces and the network interface is difficult in general. In ...

Keywords: Clusters, parallel-programming systems, system area networks

17 The case for SRAM main memory

Philip Machanick

December 1996 **ACM SIGARCH Computer Architecture News**, Volume 24 Issue 5

Full text available:  pdf(549.25 KB)

Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

The growing CPU-memory gap is resulting in increasingly large cache sizes. As cache sizes increas becomes less of a win. At the same time, since costs of going to DRAM increase, it becomes more to pin critical data in the cache---a problem if a cache is direct-mapped or has a low degree of ass Something else which is a problem for caches of low associativity is reducing misses by using a be policy. This paper proposes that L2 cache s ...

18 Resource partitioning in general purpose operating systems: experimental results in Window

D. G. Waddington, D. Hutchison

October 1999 **ACM SIGOPS Operating Systems Review**, Volume 33 Issue 4

Full text available:  pdf(1.56 MB)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

The principal role of the operating system is that of resource management. Its task is to present a services to the applications and users it supports. Traditionally, general-purpose operating system Windows NT, federate resource sharing in a fair manner, with the predominant goal of efficient res a result the chosen scheduling algorithms are not suited to applications that have stringent Quality and resource management require ...


19 Application performance and flexibility on exokernel systems

M. Frans Kaashoek, Dawson R. Engler, Gregory R. Ganger, Héctor M. Briceño, Russell Hunt, David M. Pinckney, Robert Grimm, John Jannotti, Kenneth Mackenzie

20 Load-sensitive routing of long-lived IP flows

Anees Shaikh, Jennifer Rexford, Kang G. Shin

August 1999 **ACM SIGCOMM Computer Communication Review , Proceedings of the conference**
Applications, technologies, architectures, and protocols for computer commun
Issue 4

Full text available:  pdf(1.57 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Internet service providers face a daunting challenge in provisioning network resources, due to the Internet and wide fluctuations in the underlying traffic patterns. The ability of dynamic routing to congested links and improve application performance makes it a valuable traffic engineering tool. deployment of load-sensitive routing is hampered by the overheads imposed by link-state update selection, and signaling. Under reasonable protoc ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, I
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)